

Online Multi-Commodity Flow with High Demands

Guy Even*

Moti Medina*

April 24, 2012

Abstract

This paper deals with the problem of computing, in an online fashion, a maximum benefit multi-commodity flow (ONMCF), where the flow demands may be bigger than the edge capacities of the network.

We present an online, deterministic, centralized, all-or-nothing, bi-criteria algorithm. The competitive ratio of the algorithm is constant, and the algorithm augments the capacities by at most a logarithmic factor.

The algorithm can handle two types of flow requests: (i) low demand requests that must be routed along a path, and (ii) high demand requests that may be routed using a multi-path flow.

Two extensions are discussed: requests with known durations and machine scheduling.

Keywords. Online algorithms, primal-dual scheme, multi-commodity flow.

1 Introduction

We study the problem of computing a multi-commodity flow in an online setting (ONMCF). The network is fixed and consists of n nodes and m directed edges with capacities. The adversary introduces flow requests in an online fashion.

A flow request r_j is specified by the source node s_j , the target node t_j , the demand d_j , i.e., the amount of flow that is required, and the benefit b_j , i.e., the credit that is given for a served request.

We focus on an *all-or-nothing* scenario, where a credit b_j is given only if a request r_j is fully served, otherwise, the credit is zero. Given a sequence of flow requests, the goal is to compute a multi-commodity flow (MCF) that maximizes the total benefit of *fully served* requests. Our algorithm can deal with high demands d_j . In particular, the demand d_j may be bigger than the maximum capacity.

*School of Electrical Engineering, Tel-Aviv Univ., Tel-Aviv 69978, Israel.
 {guy,medinamo}@eng.tau.ac.il

Our Contribution. We present a centralized, deterministic, all-or-nothing, non-preemptive on-line algorithm for the ONMCF problem with high demands. The algorithm is $O(1)$ -competitive. The algorithm violates edges capacities by an $O(\log n)$ factor.

We show how to extend the algorithm so that it handles two types of flow requests: (i) low demand requests that must be routed along a path, and (ii) high demand requests that may be routed using a multi-path flow.

Finally, two extensions are discussed: requests with known durations and machine scheduling.

1.1 Previous Work

Online multi-commodity flow was mostly studied in the context of *single path routing*. The *load* of an edge e in a network is the ratio between the flow that traverses e and its capacity.

Online routing was studied in two settings: (1) throughput maximization, i.e., maximizing the total benefit gained by flow requests that are served [AAP93, BN06, EMSS12], and (2) load minimization, i.e., routing all requests while minimizing the maximum load of the edges [AAF⁺97, AAPW01, BN06, BLNZ11].

In these two settings the following variants are considered: (1) permanent routing [AAF⁺97, BN06, EMSS12], (2) unknown durations [AAPW01], and (3) known durations [AAP93, EMSS12, BLNZ11].

Load Minimization. In the case of permanent routing, Aspnes et. al [AAF⁺97] designed an algorithm that augments the edge capacities by a factor of at most $O(\log n)$ w.r.t. a feasible optimal routing. Buchbinder and Naor [BN06] obtained the same result by applying a primal-dual scheme. This result can be extended to requests with high demands. The extension is based on a min-cost flow oracle that replaces the shortest path oracle.

Aspnes et. al [AAF⁺97] also showed how to use approximated oracles to allocate Steiner trees in the context of multicast virtual circuit routing. They obtained a competitive ratio of $O(\log n)$. Recently Bansal et.al [BLNZ11] extended this result to bi-criteria oracles and showed how to embed d -depth trees and cliques in the context of resource allocation in cloud computing. In the case of cliques, they required that the pairwise demands are uniform and smaller than the edge capacities. For the case of clique embedding they obtained a competitive ratio of $O(\log^3 n \cdot \log(nT))$ w.r.t. a feasible optimal solution, where T is the ratio between the maximum duration to the minimum duration of a request.

Throughput Maximization. For the case of known durations, Awerbuch et. al [AAP93] designed an $O(\log(nT))$ competitive algorithm, where T is the maximum request duration. This algorithm requires that the demands are smaller than the edge capacity by a logarithmic factor. Buchbinder and Naor [BN06, BN09b] introduced the primal-dual scheme in the online setting and designed a bi-criteria algorithm that is 1-competitive while augmenting edge capacities by a factor of $O(\log n)$ for the case of unit demands and unit benefits.

Recently Even et. al [EMSS12] showed how to apply the primal-dual scheme to embed a variety of traffic patterns in the context of Virtual Networks (VNETs). Their goal is to maximize

the profit of the served VNET requests. Some of the results in [EMSS12] require solving the ONMCF problem with high demands.

1.2 Approaches for Online MCF with High Demands

We briefly discuss the weaknesses of approaches for solving the ONMCF problem that rely directly on previous algorithms.

The algorithms in [AAP93, BN06] route each request along a single path. They require that the demand is smaller than the capacities. In order to apply these methods one should augment the capacities in advance so that the requested demand is bounded by the bottleneck along each path from the source to the destination. This augmentation might be polynomial compared to the logarithmic augmentation requirement by our algorithm.

Another option is to split the requests into subrequests of small demand so that each subdemand is smaller than the minimum capacity. After that, a single-path online routing algorithm [AAP93, BN06] can be used to route each of these subrequests. In this case, some of the subrequests might be rejected, hence violating our all-or-nothing requirement.

We define the *granularity* of a flow as the smallest positive flow along an edge in the network. Let ε denote the granularity of a flow. One can formulate the multi-commodity problem as a packing linear problem and apply the methods in [BN09b, BN09a]. The edge capacity augmentation of these algorithms depends on the $\log(\frac{1}{\varepsilon})$, which might be unbounded. For example, consider the following network: (1) The set of nodes is $V = \{u, v\}$, (2) there are two unit capacity parallel edges (u, v) . Consider a request with demand $d_j = 1 + \varepsilon$, for $\varepsilon < 1$. If the flow oracle computes an all-or-nothing flow that routes flow of size 1 on one edge and ε on the other, then the granularity is ε .

In order to solve this granularity problem, one can apply [BN09b, BN09a] and apply randomized rounding to obtain an all-or-nothing solution with unit granularity. Even in the unit-demand case, this technique increases the competitive ratio from $O(1)$ to $O(\log n)$ while the edge capacity augmentation is $O(\log n)$. Our result shows that an $O(1)$ -competitive ratio is achievable.

1.3 Techniques

Our algorithm is based on the *online primal-dual scheme*. The online primal-dual scheme by Buchbinder and Naor [BN09a, BN09b, BN06] invokes an “min-weight” path oracle. The oracles considered in [BN09a, BN09b, BN06] are either exact oracles or approximate oracles. Bansal et. al [BLNZ11] use bi-criteria oracles. Namely, the oracles they considered are approximated and augment the edge capacities. We need tri-criteria oracles.

We extend the online primal-dual scheme so it supports *tri-criteria* oracles. In the context of MCF, the oracles compute min-cost flow. The three criteria of these oracles are: (1) the approximation ratio, (2) the capacity augmentation of the edges, and (3) the granularity of the computed flow.

Multiple criteria oracles were studied by Kolliopoulos and Young [KY05]. They presented bi-criteria approximation algorithms for covering and packing integer programs. Their algorithm

finds an approximate solution while violating the packing constraints. The granularity property is used in [KY05] to mitigate this violation.

2 Problem Definition

Online multi-commodity flow (ONMCF) is defined as follows.

The Network. Let $G = (V, E)$ denote a directed graph, where V is the set of nodes and E is the set of directed edges of the network. Let $n \triangleq |V|$, and $m \triangleq |E|$. Each edge $e \in E$ has a capacity $c_e \geq 1$.

The Input. The online input is a sequence of requests σ , i.e., $\sigma = \{r_j\}_{j \in \mathbb{N}^+}$. Each flow request is a 4-tuple $r_j = (s_j, t_j, d_j, b_j)$. Let $s_j, t_j \in V$ denote, respectively, the source node and the target node of the j th request. Let $d_j \geq 1$ denote the flow demand for the j th request. Let $b_j \geq 1$ denote the benefit for the j th request. We consider an online setting, namely, the requests arrive one-by-one, and no information is known about a request r_j before its arrival.

The Output. The output is a multi-commodity flow $F = (f_1, f_2, \dots)$. For each request r_j , f_j is a flow from s_j to t_j .

Terminology. Let $|f_j|$ denote the amount of flow of f_j . Let $f_j(e)$ denote the j th flow along the edge $e \in E$. Finally, for every $e \in E$, $F^{(j)}(e) \triangleq \sum_{k=1}^j f_k(e)$, that is, the accumulated flow along an edge e after request r_j is processed. We say that an MCF $F = (f_1, f_2, \dots)$ *fully serves* a request r_j if $|f_j| = d_j$. We say that an MCF F *rejects* a request r_j if $|f_j| = 0$. We say that an MCF is *all-or-nothing* if each request is either fully served or rejected. An all-or-nothing MCF is credited b_j for each fully served request r_j . We say that an online MCF (ONMCF) algorithm is *monotone* if flow is never retracted. We say that an online MCF (ONMCF) algorithm is *preemptive* if the flow f_j of a fully served request r_j is retracted entirely, i.e., $|f_j| = 0$. A monotone ONMCF algorithm is, in particular, *non-preemptive*.

The Objective. The goal is to compute an all-or-nothing ONMCF that maximizes the total benefit of the served requests.

2.1 The Main Result

We present an online algorithm for the ONMCF problem that satisfies the following properties:

1. The algorithm is centralized and deterministic.
2. There is no limitation on demands. In particular, $\min_j d_j$ may exceed $\max_e c_e$.
3. The algorithm is all-or-nothing.

4. The online algorithm ALG competes with an all-or-nothing offline optimal algorithm.
5. The algorithm is $(1 + \delta)$ -competitive, for a constant $\delta \in (0, 1]$.
6. The algorithm violates the capacity constraints by an $O(\log n)$ factor.
7. The algorithm is non-preemptive and monotone.

For a vector $x = (x_1, \dots, x_k)$, let $x_{\min} \triangleq \min_i x_i$. Similarly, $x_{\max} \triangleq \max_i x_i$. The main result is formalized in Theorem 1.

Theorem 1 (Main Result). *Let γ denote a constant. Assume that:*

1. $1 \leq b_{\min} \leq b_{\max} \leq O(n^\gamma)$,
2. $1 \leq c_{\min} \leq c_{\max} \leq O(n^\gamma)$,
3. $1 \leq d_{\min}$.

Then, Algorithm 1 is a non-preemptive, monotone, online algorithm for the ONMCF problem that computes an all-or-nothing multi-commodity flow that is $(O(1), O(\log n))$ -competitive.

3 Online Packing and Covering Formulation

In this section we present a sequence of packing linear programs (LPs) that correspond to the ONMCF problem. We also present covering linear programs. We refer to the covering programs as the primal LPs and to the packing programs as the dual LPs.

3.1 Flow Polytopes

We define polytopes of flows that correspond to the requests $\{r_j\}_{j \in \mathbb{N}^+}$ as follows.

Definition 1. *For every $r_k = (s_k, t_k, d_k, b_k)$, let $\Pi_k(\mu)$ denote the polytope of unit flows f from s_k to t_k in G that satisfy: $\forall e \in E : f(e) \leq \mu \cdot \frac{c_e}{d_k}$.*

We refer to $\Pi_k(1)$ simply by Π_k . Let $V(\Pi_k(\mu))$ denote the set of extreme points of $\Pi_k(\mu)$.

Definition 2. *We say that request r_k is μ -feasible if $\Pi_k(\mu) \neq \emptyset$. We say that request r_k is feasible if $\Pi_k \neq \emptyset$.*

Note that a request r_j is μ -feasible if and only if the capacity of the minimum cut that separates s_k from t_k is at least $\frac{d_k}{\mu}$. In particular, a request r_j may be feasible even if $d_j > \max_e c_e$.

P-LP(j) :

$$\begin{aligned} \min \quad & \sum_{k=1}^j d_k \cdot z_k + \sum_{e \in E} c_e \cdot x_e \quad \text{s.t.} \\ \forall k \in [1, j] \quad \forall f \in V(\Pi_k) : \quad & z_k + \sum_{e \in E} x_e \cdot f(e) \geq \frac{b_k}{d_k} \\ & x, z \geq \vec{0} \end{aligned}$$

(I)

D-LP(j) :

$$\begin{aligned} \max \quad & \sum_{k=1}^j \sum_{f \in V(\Pi_k)} \frac{b_k}{d_k} \cdot y_f \quad \text{s.t.} \\ \forall e \in E : \quad & \sum_{k=1}^j \sum_{f \in V(\Pi_k)} f(e) \cdot y_f \leq c_e \quad (\text{Capacity Constraints.}) \\ \forall k \in [1, j] : \quad & \sum_{f \in V(\Pi_k)} y_f \leq d_k \quad (\text{Demand Constraints.}) \\ & y \geq \vec{0} \end{aligned}$$

(II)

Figure 1: (I) The primal LP P-LP(j). (II) The dual LP D-LP(j).

3.2 Packing and Covering Formulation

For every prefix of requests $\{r_k\}_{k=1}^j$ we define a primal linear program P-LP(j) and a dual linear program D-LP(j). The LP's appear in Figure 1.

The packing program D-LP(j) has a variable y_f for every flow $f \in \bigcup_k V(\Pi_k)$ and two types of constraints: *demand constraints* and *capacity constraints*. The capacity constraints require that the load on every edge e is at most c_e . The demand constraints require that the conical combination of unit flows in $V(\Pi_k)$ is a flow of size at most d_k .

The covering program P-LP(j) has a variable x_e for every edge $e \in E$, and a variable z_k for every request r_k , where $k \leq j$. It is useful to view x_e as the cost of a unit flow along e .

4 The Online Algorithm ALG

In this section we present the online algorithm ALG.

4.1 Preliminaries

The algorithm maintains the following variables: (1) For every edge e the primal variable x_e , (2) for every request r_j the primal variable z_j , and (3) the multi-commodity flow F . The primal variables x, z are initialized to zero. The MCF F is initialized to zero as well.

Notation. Let $x_e^{(j)}$ denote the value of the primal variable x_e after request r_j is processed by ALG. For every request r_j , let $cost_j(f)$ denote the x -cost of a flow f , formally:

$$cost_j(f) \triangleq \sum_e x_e^{(j-1)} \cdot f(e).$$

For every flow f , let $w(f)$ denote the sum of the flows along the edges, formally:

$$w(f) \triangleq \sum_e f(e).$$

Let $F^{(k)}$ denote the MCF F after request r_k is processed. Let $benefit_j(F)$ denote the benefit of MCF F after request r_j is processed, formally:

$$benefit_j(F) \triangleq \sum \{b_i \mid i \leq j, r_i \text{ is fully served by } F^{(j)}\}.$$

Let $value_j(x, z)$ denote the objective function's value of P-LP(j) for a given x and z , formally:

$$value_j(x, z) \triangleq \sum_{k=1}^j d_k \cdot z_k + \sum_{e \in E} c_e \cdot x_e^{(j)}.$$

Let F^* denote an all-or-nothing offline optimal MCF w.r.t input sequence $\sigma = \{r_j\}_j$.

Definition 3. An MCF $F = (f_1, f_2, \dots)$ is (α, β) -competitive with respect to a sequence $\{r_j\}_j$ of requests if for every j :

- (i) F is α -competitive: $\text{benefit}_j(F) \geq \frac{1}{\alpha} \cdot \text{benefit}_j(F^*)$.
- (ii) F is β -feasible: for every $e \in E$, $F^{(j)}(e) \leq \beta \cdot c_e$.

Definition 4. An MCF $F = (f_1, f_2, \dots)$ is all-or-nothing if each request r_j is either fully served by F or it is rejected by F (i.e., $|f_j| \in \{0, d_j\}$).

4.2 Description

Upon arrival of a request r_j , if the request is not feasible, then the algorithm rejects it upfront. Otherwise, if the request is feasible, then ALG invokes a tri-criteria oracle. The oracle returns a unit-flow f_j for r_j .

If the cost of the oracles's flow is “small enough”, then the request is accepted as follows: (1) the flow F is updated by adding the oracle's unit-flow f_j times the required demand d_j , (2) the primal variables x_e , for every edge e that the flow f_j traverses, are updated.

If the flow is “too expensive”, then the request is rejected and no updates are made to the primal variables and to the MCF F .

The listing of the online algorithm ALG appears in Algorithm 1.

4.3 The Oracle

The oracle description is as follows:

- (i) **Input:** Request r_j , edge capacities $\frac{c_e}{d_j}$, and edge costs $x^{(j-1)} : E \rightarrow \mathbb{R}^{\geq 0}$.
- (ii) **Output:** A unit-flow f from s_j to t_j .

Let MIN-COST_j denote the min-cost flow in Π_j w.r.t. the costs x_e , formally:

$$\text{MIN-COST}_j \triangleq \arg \min \{ \text{cost}_j(f) : f \in \Pi_j \}.$$

Note that: (1) MIN-COST_j is well defined because $\Pi_j \neq \emptyset$, and (2) the edge capacities in Π_j are $\frac{c_e}{d_j}$.

The oracles in our context are tri-criteria, as formalized in the following definition.

Definition 5 (Oracle Criteria). We say that an oracle is $(\lambda, \mu, \varepsilon)$ -criteria, if the oracle outputs a flow f that satisfies the following properties:

- (i) (λ -**Approximation.**) $\text{cost}_j(f) \leq \lambda \cdot \text{cost}_j(\text{MIN-COST}_j)$.
- (ii) (μ -**Augmentation.**) $f \in \Pi_j(\mu)$.
- (iii) (ε -**Granular.**) $f(e) > 0 \Rightarrow f(e) \geq \varepsilon$.

Algorithm 1 ALG: Online multi-commodity flow algorithm. The algorithm receives a sequence of requests and outputs a multi-commodity flow F .

Initialize: $z \leftarrow 0, x \leftarrow 0, F \leftarrow 0$.

Upon arrival of request $r_j = (s_j, t_j, d_j, b_j)$, for $j \geq 1$:

- 1) If r_j is not feasible (i.e., $\Pi_j = \emptyset$), then **reject** r_j and skip the remaining lines.
- 2) $f_j \leftarrow \text{oracle}(x, r_j)$ {The oracle is a $(\lambda, \mu, \varepsilon)$ -criteria.}
- 3) If $d_j \cdot \text{cost}_j(f_j) < \lambda \cdot b_j$,
- 4) then **accept** r_j
 - 5) $F \leftarrow F + d_j \cdot f_j$ {Updating the multi-commodity flow.}
 - 6) $z_j \leftarrow \frac{b_j}{d_j} - \frac{\text{cost}_j(f_j)}{\max\{\lambda, \mu\}}$
 - 7) $\forall e : f_j(e) > 0$:

$$L_j(e) \triangleq \frac{d_j \cdot f_j(e)}{\max\{\lambda, \mu\} \cdot c_e}$$

$$x_e \leftarrow x_e \cdot 2^{L_j(e)} + \frac{1}{d_j \cdot w(f_j)} \cdot (2^{L_j(e)} - 1)$$

- 8) Else **reject** r_j
-

4.3.1 A Tri-criteria Oracle for Minimum Cost Flow

The oracle's listing is as follows.

The Oracle Outline.

1. Let $f \leftarrow \text{MIN-COST}_j$.
2. Decompose f to at most m flow paths $\{f_1, \dots, f_m\}$.
3. Remove each flow path f_ℓ such that $|f_\ell| < \frac{1}{2m^2}$.
4. Let g denote the removed flow from f .
5. Scale every remaining flow path f_ℓ (i.e., $|f_\ell| \geq \frac{1}{2m^2}$) as follows:

$$f_\ell \leftarrow f_\ell \cdot \left(1 + \frac{|g|}{|f| - |g|}\right)$$

The proof of the following lemma appears in Appendix A.

Lemma 1. *The oracle is $(2, 2, \frac{1}{2m^2})$ -criteria algorithm.*

Lemma 1 justifies using the following parameters: $\lambda = \mu = 2$, and $\varepsilon = \frac{1}{2m^2}$.

5 Analysis

The following observation is proved by the fact that $\frac{1}{z} \cdot (2^z - 1)$ is monotone increasing for $z > 0$.

Observation 1. *Let $c \in \mathbb{R}^{>0}$, then*

$$\forall x \in [0, c] : c \cdot (2^{x/c} - 1) \leq x.$$

Observation 2. *If $L_j(e) \leq 1$, then*

$$(2^{L_j(e)} - 1) \cdot c_e \leq \frac{1}{\max\{\lambda, \mu\}} \cdot d_j \cdot f_j(e).$$

Proof. By Observation 1 and since $L_j(e) \leq 1$, it follows that

$$(2^{L_j(e)} - 1) \cdot \max\{\lambda, \mu\} \cdot c_e \leq d_j \cdot f_j(e),$$

and the observation follows. □

Notation. Let

$$\alpha \triangleq 1 + \frac{1}{\max\{\lambda, \mu\}} \leq 2,$$

$$\beta \triangleq \max\{\lambda, \mu\} \cdot \log_2 \left(1 + m^2 \cdot \frac{3 \cdot \lambda \cdot c_{\max} \cdot b_{\max}}{\varepsilon} \right).$$

In the following theorem we prove that ALG is an all-or-nothing (α, β) -competitive, non-preemptive and monotone online algorithm.

Theorem 2. *Assume that:*

1. $b_{\min}, c_{\min}, d_{\min} \geq 1$.
2. *The oracle is $(\lambda, \mu, \varepsilon)$ -criteria.*

Then ALG is non-preemptive, monotone, online algorithm for the ONMCF problem that computes an all-or-nothing multi-commodity flow that is (α, β) -competitive.

Proof. The ALG algorithm rejects upfront requests that are not feasible. These requests are also rejected by F^* , hence it suffices to prove (α, β) -competitiveness w.r.t fractional offline optimal algorithm over the feasible requests. We now prove α -competitiveness and β -feasibility.

α -competitiveness. First, we prove α -competitiveness. Let $\Delta_j P \triangleq \text{value}_j(x, z) - \text{value}_{j-1}(x, z)$, and $\Delta_j F \triangleq \text{benefit}_j(F) - \text{benefit}_{j-1}(F)$. We begin by proving that $\Delta_j P \leq \alpha \cdot \Delta_j F$ for every request r_j .

Recall that $x_e^{(j)}$ denotes the value of the primal variable x_e after r_j is processed. If r_j is rejected then $\Delta_j P = \Delta_j F = 0$ and the claim holds. If r_j is accepted, then $\Delta_j F = b_j$ and $\Delta_j P = \sum_e (x_e^{(j)} - x_e^{(j-1)}) \cdot c_e + d_j \cdot z_j$. Let f_j denote the output of the oracle when dealing with request r_j , i.e., $f_j \leftarrow \text{oracle}(x^{(j-1)}, r_j)$. Indeed,

$$\begin{aligned} \sum_e (x_e^{(j)} - x_e^{(j-1)}) \cdot c_e &= \sum_e \left[x_e^{(j-1)} \cdot (2^{L_j(e)} - 1) + \frac{1}{d_j \cdot w(f_j)} \cdot (2^{L_j(e)} - 1) \right] \cdot c_e \\ &= \sum_e \left(x_e^{(j-1)} + \frac{1}{d_j \cdot w(f_j)} \right) \cdot (2^{L_j(e)} - 1) \cdot c_e \\ &\leq \sum_e \left(x_e^{(j-1)} + \frac{1}{d_j \cdot w(f_j)} \right) \cdot \frac{d_j \cdot f_j(e)}{\max\{\lambda, \mu\}} \\ &= \frac{d_j \cdot \text{cost}_j(f_j)}{\max\{\lambda, \mu\}} + \frac{1}{\max\{\lambda, \mu\}}, \end{aligned} \tag{1}$$

where the third inequality holds since the oracle is μ -augmented and by Observation 2. Hence,

Equation 1 and Step 6 of ALG imply that:

$$\begin{aligned}
\Delta_j P &\leq \frac{d_j \cdot \text{cost}_j(f_j)}{\max\{\lambda, \mu\}} + \frac{1}{\max\{\lambda, \mu\}} + d_j \cdot z_j \\
&= \frac{d_j \cdot \text{cost}_j(f_j)}{\max\{\lambda, \mu\}} + \frac{1}{\max\{\lambda, \mu\}} + d_j \cdot \left(\frac{b_j}{d_j} - \frac{\text{cost}_j(f_j)}{\max\{\lambda, \mu\}} \right) \\
&= \frac{1}{\max\{\lambda, \mu\}} + b_j \leq \alpha \cdot b_j,
\end{aligned}$$

where the last inequality holds since $b_j \geq 1$.

Since $\Delta_j F = b_j$ it follows that

$$\Delta_j P \leq \alpha \cdot \Delta_j F, \quad (2)$$

as required.

Initially, the primal variables and the flow F equal zero. Hence, Equation 2 implies that:

$$\text{value}_j(x, z) \leq \alpha \cdot \text{benefit}_j(F). \quad (3)$$

We now prove that, the primal variables $\{x_e^{(j)}\}_e \cup \{z_i\}_{i \leq j}$ constitute a feasible solution for P-LP(j):

1. If r_j is rejected, then $\text{cost}_j(f_j) \geq \lambda \cdot \frac{b_j}{d_j}$. Since the oracle is λ -approximate it follows that for every $f' \in V(\Pi_j)$:

$$\text{cost}_j(f') \geq \text{cost}_j(\text{MIN-COST}_j) \geq \text{cost}_j(f_j) / \lambda \geq \frac{b_j}{d_j}.$$

It follows that the primal constraints are satisfied in this case.

2. If r_j is accepted, then $\text{cost}_j(f_j) < \lambda \cdot \frac{b_j}{d_j}$. Since $z_j = \frac{b_j}{d_j} - \frac{\text{cost}_j(f_j)}{\max\{\lambda, \mu\}}$ it follows that for every $f' \in V(\Pi_j)$:

$$z_j + \text{cost}_j(f') \geq \frac{b_j}{d_j} - \frac{\text{cost}_j(f_j)}{\max\{\lambda, \mu\}} + \frac{\text{cost}_j(f_j)}{\lambda} \geq \frac{b_j}{d_j}.$$

We conclude that the primal constraints are satisfied in this case as well.

The first j flows of the optimal offline multi-commodity flow F^* are clearly a feasible solution to D-LP(j). The value of this solution equals $\text{benefit}_j(F^*)$. Since the primal variables constitute a feasible primal solution, weak duality implies that:

$$\text{benefit}_j(F^*) \leq \text{value}_j(x, z).$$

Hence, by Equation 3, it follows that:

$$\text{benefit}_j(F) \geq \frac{1}{\alpha} \cdot \text{benefit}_j(F^*),$$

which proves that ALG is α -competitive.

β -feasibility. We now prove β -feasibility, i.e., for every r_i and for every $e \in E$, $F^{(i)}(e) \leq \beta \cdot c_e$.

We prove a lower bound and an upper bound on x_e in the next two lemmas. Let r_j denote the index of the last request. Let

$$W \triangleq \max\{d_k \cdot w(f_k) : 0 \leq k \leq j\}.$$

Lemma 2. For every edge e ,

$$x_e \geq \frac{1}{W} \cdot \left(2^{F(e)/\max\{\lambda, \mu\} \cdot c_e} - 1\right).$$

Proof. Recall that $x^{(k)}$ (resp. $F^{(k)}$) denote the value of x (resp. F) after request r_k is processed. We prove by induction on $k \leq j$ that

$$x_e^{(k)} \geq \frac{1}{W} \cdot \left(2^{F^{(k)}(e)/\max\{\lambda, \mu\} \cdot c_e} - 1\right). \quad (4)$$

The induction basis, for $k = 0$, holds because both sides equal zero.

Induction step: Note that if r_k is rejected, then both sides of Equation 4 remain unchanged, and hence Equation 4 holds by the induction hypothesis. We now consider the case that r_k is accepted.

The update rule in Step 7 of ALG implies that

$$\begin{aligned} x_e^{(k)} &= x_e^{(k-1)} \cdot 2^{\frac{d_k \cdot f_k(e)}{\max\{\lambda, \mu\} \cdot c_e}} + \frac{1}{d_j \cdot w(f_k)} \cdot \left(2^{\frac{d_k \cdot f_k(e)}{\max\{\lambda, \mu\} \cdot c_e}} - 1\right) \\ &\geq \frac{1}{W} \cdot \left(2^{\frac{F^{(k-1)}(e)}{\max\{\lambda, \mu\} \cdot c_e}} - 1\right) \cdot 2^{\frac{d_k \cdot f_k(e)}{\max\{\lambda, \mu\} \cdot c_e}} + \frac{1}{d_j \cdot w(f_k)} \cdot \left(2^{\frac{d_k \cdot f_k(e)}{\max\{\lambda, \mu\} \cdot c_e}} - 1\right) \\ &\geq \frac{1}{W} \cdot \left(2^{\frac{F^{(k)}(e)}{\max\{\lambda, \mu\} \cdot c_e}} - 1\right). \end{aligned}$$

The lemma follows. □

Lemma 3. For every accepted request r_k , if $f_k(e) > 0$, then

$$x_e^{(k)} \leq \frac{3 \cdot \lambda \cdot b_k}{\varepsilon \cdot d_k}.$$

Proof. Since r_k is accepted, we have $d_k \cdot \text{cost}_k(f_k) < \lambda \cdot b_k$. By ε -granularity of the oracle, $\text{cost}_k(f_k) \geq x_e^{(k-1)} \cdot \varepsilon$. It follows that $x_e^{(k-1)} \leq \frac{\lambda \cdot b_k}{d_k \cdot \varepsilon}$. By the update rule for x_e , we have:

$$x_e^{(k)} \leq \frac{\lambda \cdot b_k}{d_k \cdot \varepsilon} \cdot 2^{L_k(e)} + \frac{1}{d_k \cdot w(f_k)} \cdot (2^{L_k(e)} - 1).$$

Since the oracle is μ -augmented, $L_k(e) \leq 1$. In addition, since the oracle is ε -granular, $w(f_j) \geq \varepsilon$.

$$\begin{aligned} x_e^{(k)} &\leq \frac{\lambda \cdot b_k}{d_k \cdot \varepsilon} \cdot 2 + \frac{1}{d_k \cdot \varepsilon} \cdot (2 - 1) \\ &\leq \frac{3 \cdot \lambda \cdot b_k}{d_k \cdot \varepsilon}. \end{aligned}$$

□

Lemma 2 and Lemma 3 imply that:

$$\frac{1}{W} \cdot \left(2^{F^{(k)}(e) / \max\{\lambda, \mu\} \cdot c_e} - 1 \right) \leq \max_{k \leq j} \frac{3 \cdot \lambda \cdot b_k}{\varepsilon \cdot d_k}.$$

Hence,

$$F^{(k)}(e) \leq c_e \cdot \max\{\lambda, \mu\} \cdot \log_2 \left(1 + W \cdot \max_k \frac{3 \cdot \lambda \cdot b_k}{\varepsilon \cdot d_k} \right). \quad (5)$$

Since (i) $W \leq m \cdot d_{\max}$, (ii) $d_{\max} \leq m \cdot c_{\max}$, and (iii) $d_{\min} \geq 1$, it follows that,

$$F^{(k)}(e) \leq \beta \cdot c_e,$$

for every k , as required. \square

This concludes the proof of Theorem 2. Theorem 1 follows directly from Theorem 2 and Lemma 1.

Remark 1. Let bpb_k denote the benefit-per-bit of request r_k , i.e., $bpb_k \triangleq \frac{b_k}{d_k}$. Let $bpb_{\max} \triangleq \max_k bpb_k$. Instead of β , the augmentation can be also bounded by:

$$\max\{\lambda, \mu\} \cdot \log_2 \left(1 + W \cdot \frac{3 \cdot \lambda}{\varepsilon} \cdot bpb_{\max} \right).$$

6 Mixed Demands

One may consider a mixed case of low and high demands. A flow request with high demand has to be split into multiple paths. Splitting a stream of packets along multiple paths should avoided, if possible, because it complicates implementation in nodes where flow is split, may cause packets to arrive out-of-order, etc. Thus, one may require not to split requests with low demand. Formally, a request has low demand if $d_j \leq c_{\min}$; otherwise, it has a high demand.

An online algorithm for mixed demands can be obtained by employing two oracles: (1) A tri-criteria oracle for the high demands. This oracle may serve a flow request by multiple paths. (2) An exact (shortest path) oracle for low demands. This oracle must serve a flow request by a single path.

Theorem 3. *There exists a non-preemptive, monotone, online algorithm for the ONMCF problem with mixed demands that computes an all-or-nothing multicommodity flow that is $(O(1), O(\log n))$ -competitive.*

Proof sketch. The proof is based on the feasibility of the primal LP and on the bounded gap between $\Delta_j F$ and $\Delta_j P$. These two invariants are maintained regardless of the oracle that is invoked. The proof for the case of small demands appears in [BN09a]. The augmentation of the capacities are determined by the oracle with the “worst” parameters. Because the exact oracle is $(1, 1, 1)$ -criteria, it is also (λ, μ, ϵ) -criteria. Thus, the augmentation factor β is determined by the approximate oracle. \square

7 Further Extensions

Requests with known durations. The algorithm can be extended to deal with flow requests with known durations. For the sake of simplicity, the flow requests in this paper are permanent, namely, after arrival, a request stays forever. Using previous techniques [AAP93, BN06, EMSS12], our algorithm can be adapted to deal also with the important variant of known durations. In this variant, each request, upon arrival, also has an end-time. The competitive ratio for known durations when the requests are a logarithmic fraction of the capacities is $O(\log(nT))$, where T denotes the longest duration [AAP93]. In fact, the primal-dual method in [BN06] can be extended to the case of routing requests with known durations (see [EMSS12]). Thus, for known durations, if the demands are bounded by the minimum capacity, then the primal-dual method yields an online algorithm, the competitive ratio of which is $(O(1), O(\log(nT)))$. One can apply a tri-criteria oracle with granularity $O(n^{-2})$, to obtain an $(O(1), O(\log(nT)))$ -competitive ratio for known durations even with high demands.

All-or-nothing machine scheduling. A simple application of our algorithm is the case of maximizing throughput in an online job all-or-nothing scheduling problem on unrelated machines. The variant in which the objective is to minimize the load was studied by Aspnes et al. [AAF⁺97]. We, on the other hand, focus on maximizing the throughput.

Jobs arrive online, and may be assigned to multiple machines immediately upon arrival. Moreover, a job may require specific subset of machines, i.e., restricted assignment. The increase in the load of a machine when a job is assigned to it is a function of the machine and the fraction of the job that is assigned to it. Formally, Let $\tau_j(e) \in [0, 1]$ denote the “speed up” of machine e when processing job j , that is, one unit of job e on machine j incurs a an additional load of $\tau_j(e)$ on machine e . The reduction is to network of m parallel edges, one edge per machine. The capacity of each edge equals the capacity of the corresponding machine.

Large jobs need to be assigned to multiple machines, while small jobs may be assigned to a single machine (as in [AAF⁺97]). In this case our algorithm is $(O(1), O(\frac{\log m}{\min_{j,e} \tau_j(e)}))$ -competitive, where m is the number of machines.

References

- [AAF⁺97] J. Aspnes, Y. Azar, A. Fiat, S. Plotkin, and O. Waarts. On-line routing of virtual circuits with applications to load balancing and machine scheduling. *Journal of the ACM (JACM)*, 44(3):486–504, 1997.
- [AAP93] B. Awerbuch, Y. Azar, and S. Plotkin. Throughput-competitive on-line routing. In *FOCS '93: Proceedings of the 1993 IEEE 34th Annual Foundations of Computer Science*, pages 32–40, Washington, DC, USA, 1993. IEEE Computer Society.
- [AAPW01] B. Awerbuch, Y. Azar, S. Plotkin, and O. Waarts. Competitive routing of virtual circuits with unknown duration. *Journal of Computer and System Sciences*, 62(3):385–397, 2001.

- [BLNZ11] Nikhil Bansal, Kang-Won Lee, Viswanath Nagarajan, and Murtaza Zafer. Minimum congestion mapping in a cloud. In *PODC*, pages 267–276, 2011.
- [BN06] Niv Buchbinder and Joseph (Seffi) Naor. Improved bounds for online routing and packing via a primal-dual approach. *Foundations of Computer Science, Annual IEEE Symposium on*, 0:293–304, 2006.
- [BN09a] Niv Buchbinder and Joseph (Seffi) Naor. The design of competitive online algorithms via a primal-dual approach. *Foundations and Trends in Theoretical Computer Science*, 3(2-3):99–263, 2009.
- [BN09b] Niv Buchbinder and Joseph (Seffi) Naor. Online primal-dual algorithms for covering and packing. *Math. Oper. Res.*, 34(2):270–286, 2009.
- [EMSS12] Guy Even, Moti Medina, Gregor Schaffrath, and Stefan Schmid. Competitive and deterministic embeddings of virtual networks. In *ICDCN*, pages 106–121, 2012. (see also Corr abs/1101.5221, Jan. 2011).
- [KY05] S.G. Kolliopoulos and N.E. Young. Approximation algorithms for covering/packing integer programs. *Journal of Computer and System Sciences*, 71(4):495–505, 2005.

A Proof of Lemma 1

In this section we prove the following lemma.

Lemma 1 *The oracle is $(2, 2, \frac{1}{2m^2})$ -criteria algorithm.*

Proof. Throughout this proof we refer to flow path that are not removed, simply by ‘flow paths’.

First, we prove that the oracle outputs a unit flow. For every flow path f_ℓ such that $|f_\ell| \geq \frac{1}{2m^2}$, let $f_\ell^{(s)}$ denote the scaled flow along it. The sum of the flows, along the scaled flow paths equals:

$$\sum_{\ell} f_\ell^{(s)} = \sum_{\ell} f_\ell \cdot \left(1 + \frac{|g|}{|f| - |g|}\right) = |f| - |g| + |g| = |f|.$$

Hence, the oracle outputs a unit flow as required.

The oracle is $\frac{1}{2m^2}$ -granular by construction.

We prove that the oracle is 2-augmented. Note that $|g| \leq \frac{m}{2m^2} = \frac{1}{2m}$. Hence,

$$\frac{|g|}{|f| - |g|} \leq \frac{\frac{1}{2m}}{1 - \frac{1}{2m}} = \frac{1}{2m - 1}.$$

It follows that the flow along every edge is augmented by at most

$$\left(1 + \frac{|g|}{|f| - |g|}\right) \leq \left(1 + \frac{1}{2m - 1}\right) < 2.$$

Hence, $f \in \Pi_j(2)$, as required.

Moreover, the flow along every scaled flow path $f_\ell^{(s)}$ satisfies for every $e \in E$:

$$f_\ell^{(s)}(e) \leq f_\ell(e) \cdot \left(1 + \frac{1}{2m-1}\right),$$

Hence,

$$\text{cost}_j(f_\ell^{(s)}) \leq \text{cost}_j(f_\ell) \cdot \left(1 + \frac{1}{2m-1}\right),$$

which proves 2-approximation of the oracle. □